

A Multibody Formulation For Three Dimensional Brick Finite Element Based Parallel and Scalable Rotor Dynamic Analysis

Anubhav Datta
ELORET Corporation
AFDD at Ames Research Center
Moffett Field, CA 94035

Wayne Johnson
Aeromechanics Branch
NASA Ames Research Center
Moffett Field, CA 94035

ABSTRACT

This paper presents a unified formulation for a 3-dimensional finite element based non-linear multibody analysis for helicopter rotors. Special multibody brick elements are developed which can be used to embed arbitrary joint rotations within a 3-dimensional structure. A multi-level iterative substructuring algorithm, that is parallel and scalable, is redesigned to accommodate multibody components in a manner that maintains its underlying numerical scalability. The brick multibody formulation is then used to study the impact of non-linear 3-dimensional hub end effects in rotors that are not modeled by current generation beam based models. The constraints that arise from a physical 3-D connection of a rotor blade to a joint are shown to alter the internal stresses at its hub end and impact torsion dynamics significantly – a physics that is uniquely rotary wing in nature. Large scale structural models are then constructed for a hingeless, an articulated, and a bearingless rotor, consisting of multiple flexible components and multibody connections near the hub end, and containing up to 0.48 million degrees of freedom. The models are analyzed for scalability and timing for hover and forward flight solutions on up to 128 processors. The key conclusion is that multibody components can indeed be incorporated within a fully parallel multi-level iterative substructuring algorithm without impacting its numerical scalability. And, integrated carefully within 3-dimensional brick elements, they open new opportunities for capturing fundamental physics of 3-dimensional stress fields on rotary wing structures.

INTRODUCTION

The objective of this paper is to provide a unified formulation for a 3-dimensional (3-D) brick finite element method (FEM) based multibody dynamics analysis for helicopter rotors. Such a formulation requires an innovative method to formulate multibody joint components within a non-linear 3-D FEM analysis, and an innovative method to accommodate these components within a fully parallel and scalable solution procedure. This paper describes the development of both these methods.

This research is targeted towards the development of a high fidelity, 3-D FEM based, parallel and scalable Computational Structural Dynamics (CSD) solver for helicopter rotors. It is envisioned to be a central component of a next generation, High Performance Comput-

ing (HPC) based, high fidelity rotorcraft analysis [1]. A research effort was initiated recently by the authors in Ref. [2] towards the development of such a solver.

The current state of the art in rotorcraft dynamic analysis incorporates multibody formulations for nonlinear beams [3, 4, 5]. A vast literature exists for multibody formulations of nonlinear structural elements like beams and shells. A recent review of the status of current research on such methods can be found in Ref. [6]. Here, the intent is to devise a formulation that can be integrated with nonlinear 3-D brick finite elements. Only a limited number of studies can be found on multibody formulations of 3-D elements [7, 8, 9]. Integration of multibody dynamics will enable the 3-D rotor analysis to model realistic hub kinematics while including multiple load bearing 3-D flexible components.

The state-of-the-art in finite element analysis of helicopter blades involves a variational-asymptotic reduction of the 3-D nonlinear elasticity problem into a 2-D linear cross-section analysis and a 1-D geometrically exact beam analysis – based on Berdichevsky [10] and

Presented at the American Helicopter Society 66th Annual Forum, Phoenix, Az, May 11–13, 2010. This is a work of the U.S. Government and is not subject to copyright protection.

pioneered by Hodges et al. [11]. Aeroelastic computations are performed on the beam, followed by a recovery of the 3-D stress field. The method is efficient and accurate – except near end-edges and discontinuities for which a 3-D analysis is still needed – as long as the cross-sectional characteristic dimensions are small compared to the wavelength of deformations along the beam. Modern hingeless and bearingless rotors contain 3-D flexible components near the hub that have short aspect ratios, open sections, and end constraints, and hence cannot be treated as beams. The critical couplings that determine blade dynamics are dominated by these components. Critical stresses often occur in these same components. Moreover, the treatment of blades, depending on their advanced geometry and material anisotropy, require continuous refinements to beam modeling and analysis to accommodate new physics. The objective of the present research is to develop a 3-D FEM based rotor dynamic analysis that can model generic 3-D components and dramatically increase the scope of analysis for modern rotors.

With the emergence of rotorcraft Computational Fluid Dynamics (CFD), physics-based models containing millions of grid points carry out Reynolds Averaged Navier-Stokes (RANS) computations on hundreds of cores, routinely, in a research environment for the rotor, and even for the entire helicopter. Applications today are focused on coupling CFD with relatively simple engineering-level structural models – carried out on a single processor while the remaining processors lie idle. Assessments of the state-of-the-art in loads prediction, however, make it clear that the progress has mostly been in airloads, and much less in the accuracy of structural loads [12, 13]. The intent of this research is to explore the possibility of using 3-D FEM as the physics-based counterpart in the structures domain.

There is no question that such a capability will be powerful. First, it will enable the modeling of critical couplings that occur in hingeless and bearingless hubs with advanced flex structures. Second, it will enable the direct calculation of stresses in these critical load bearing components. Third, it will provide an equal fidelity of representation of the physics of structures and fluids, unlike the CFD/CSD simulations of today which are named so merely for the symmetry of terminology. And finally, even though this research is targeted towards HPC based analysis, it will always provide as a by-product a means (via static analysis) for extracting sectional properties with which efficient lower order beam analyses can be carried out when desired. The key questions for such a capability are, first, whether an efficient solution procedure can be found that is fully parallel and scalable. Second, whether multibody dynamics can be integrated within a non-linear 3-D brick FEM formulation. Third, whether incorporating multibody dynamics will still retain the efficiency and scalability of the procedure. The primary focus of the present research has therefore been

on answering these key questions directly.

Initial work documented in Refs. [2] and [14] demonstrated that a pure 3-D FEM based rotor dynamic analysis (i.e., without multibody dynamics) can indeed be carried out in a fully parallel and scalable manner. An advanced multi-level iterative substructuring method — the Dual-Primal Finite Element Tearing and Interconnecting (FETI-DP) method pioneered by Farhat et al. [15, 16], was used to develop and study a parallel and scalable solution of a simple 3-D rotary wing structural dynamics prototype.

The focus of the present paper is on multibody dynamics — its formulation within 3-D FEM, its impact on the physics of 3-D effects, its accommodation within a scalable iterative substructuring method, and its impact on the efficient parallel solution of large scale models of realistic rotor configurations.

Scope of Present Work

The main emphasis in this work is on the development of a special multibody dynamics formulation — one that can be integrated within 3-D non-linear brick FEM, and one that can be solved in a fully parallel and scalable manner. The constraints to be considered are holonomic in nature – adequate for most advanced rotor configurations. The formulation to be studied is restricted to an isolated rotor. Advanced modeling like rotating-non rotating interfaces, nonholonomic constraints, and friction contact, are beyond the scope of this initial work. For the purposes of solver development, and for the fundamental understanding of the requirements for advanced level modeling, simple grid generators, mergers, and partitioners are all developed in-house as part of this study. Many key elements of a comprehensive rotorcraft analysis are not considered at present: airloads, trim, and extraction of periodic dynamics are all part of future work.

The paper is organized into four sections. The first section is on the 3-D FEM multibody analysis development. It covers joint formulation, joint modeling, and physical insights into non-linear 3-D edge effects. The second section is on the development of parallel and scalable solvers for the analysis. The third section describes the key components of the analysis: geometry and grids, unique substructuring requirements with multibody dynamics, and a description of the hover and forward flight prototypes. The final section details the scalability and timings of the 3-D FEM multibody analysis for hover and forward flight calculations of large scale rotor models.

3-D FEM-MULTIBODY FORMULATION

FEM Formulation

The equations of motion are derived using generalized Hamilton's Principle governing the motion of a




Figure 1: **A 27-node isoparametric, hexahedral brick element; $4 \times 4 \times 4$ Gauss integration points.**

non-conservative system between times t_1 and t_2

$$\int_{t_1}^{t_2} (\delta U - \delta T - \delta W) dt = 0 \quad (1)$$

where δU , δT , and δW are variation in strain energy, variation in kinetic energy, and virtual work respectively. The expressions for each of these are derived in Ref. [2]. The formulation uses Green-Lagrange strains and second Piola-Kirchhoff stress measures for strain energy. The non-linear, geometrically exact implementation follows the standard Total Lagrangian based incremental approach [18, 19]. The geometrically exact formulation is a pre-requisite for the multibody formulation presented in this paper. The stress-strain material relationship is assumed to be linear.

The isoparametric, hexahedral, quadratic brick element used in this work is as shown in Fig. 1. It consists of 8 vertex nodes and 19 internal nodes – 12 edge nodes, 6 face nodes, and 1 volume node. The presence of sufficient internal nodes prevents locking. Within isoparametric elements, geometry and displacement solution are both interpolated using the same shape functions. The shape functions are expressed in element natural coordinates ξ , η , and ζ , where $-1 \leq \xi, \eta, \zeta \leq 1$. We consider 2nd order Lagrange polynomials in each direction.

$$H_a(\xi, \eta, \zeta) = L_I^n(\xi) L_J^m(\eta) L_K^p(\zeta) \quad (2)$$

where H is a shape function and a its node point index. Here $n = m = p = 2$; and I, J, K are node numbers in each direction varying as 1, 2, 3 respectively. Based on the local node ordering shown in Fig. 1(b), we have for




Figure 2: **Blade frequencies vs. normalized rotational speed for a soft in-plane hingeless rotor; collective 20° , twist -15°**




Figure 3: **Blade frequencies vs. normalized rotational speed for a fully articulated rotor; collective 20° , twist -15°**

example the shape function corresponding to node 11

$$H_{11} = L_2^2(\xi) L_3^2(\eta) L_1^2(\zeta) = \frac{1}{4} \eta \zeta (1 - \xi^2) (\eta + 1) (\zeta - 1)$$

The construction of the finite element matrices then follow as given in Ref. [2].

The brick elements are verified using the rotating frequencies of a slender beam-like geometry of rectangular cross-section, high aspect ratio, and uniform chord. The beam has dimensions of $20c \times c \times c/4$ in span, chord, and thickness directions; a uniform twist of -15° about mid-chord; and is set at a collective pitch angle of 20° . The rotation axis is along mid-chord. The material moduli are $E = 8.27 \times 10^7$ Pa and $G = 3.45 \times 10^7$ Pa ($\nu = 0.2$), density is $\rho = 192$ kg/m³, and $c = 0.0864$ m. The discretization uses $16 \times 4 \times 2$ elements denoting the number of bricks along span, chord, and thickness, respectively. The frequency plot for a hingeless blade is

shown in Fig. 2, compared with converged second-order non-linear beam element results (40 elements). The 3-D boundary conditions are zero deflections at all root nodes. The rotor speed is normalized with respect to a baseline value of 27 rad/s. The frequency plot for a fully articulated blade (5% hinge offset) is shown in Fig. 3. The articulation is assumed to have zero hinge stiffness. This boundary condition can be easily reproduced by prescribing zero deflections at a single node at the root end. Implementing non-zero hinge stiffness require multibody joint modeling in 3-D FEM. This is because, unlike beams (or structural elements), there are no rotational states in bricks (or solid elements) on which to apply rotational stiffness directly.




Figure 4: A joint formulation connecting multiple brick elements.

Joint Formulation

A joint connects and constrains the relative motion of several structural components. One component is considered the attachment side, the rest are the connection sides (Fig. 4). A joint is attached to the attachment side structure at a point A . The inputs to the joint are the motions of A . The joint variables are displacements u^J and rotations θ^J . The joint is connected to the connection side structure at an arbitrary number of points P . In the present method, these connections are realized by special purpose brick elements, called joint elements. A joint is formulated as part of these special purpose brick elements. Any number of nodes of this joint element can be connected to the joint. At a connection node, the brick DOFs are reduced to joint DOFs by an exact transformation. The method is kinematically exact because the geometry of the brick elements are exact. The motion of a connection node P relative frame F is related to the motion of attachment node A relative frame F via joint

motion as follows.

$$\begin{aligned} u^P &= u^A + u^J + C^{FJ} l \\ \Delta u^P &= \Delta u^A + \Delta u^J + G \Delta \theta^J \end{aligned} \quad (3)$$

The joint displacements are the DOFs $u^J = [u, v, w]^J$. The joint rotation matrix C^{FJ} (from frame J to F) is parameterized in this study by Euler angle DOFs $\theta^J = [\beta, \zeta, \theta]^J$. l is the undeformed position of a connection node relative to joint in joint frame. $G = \partial C^{FJ} l / \partial \theta^J$. The velocity and accelerations follow from above. The resultant mass, stiffness, damping are non-linear functions of the joint variables.




Figure 5: Large rotation about a flap joint.




Figure 6: Large rotation about a pitch joint.

The formulation is verified using static deformations on a stiff beam-like structure for which exact joint rotations can be calculated. Consider a uniform cantilevered beam of dimensions $20c \times c \times c$, $c = 0.0864\text{m}$, artificially stiffened $E = 8.27 \times 10^{10} \text{ N/m}^2$, and with a top surface pressure of $1 \times 10^3 \text{ N/m}^2$. The beam is first discretized into $8 \times 2 \times 2$ brick elements and then subdivided into two parts with a joint connection in between. The joint translations are restrained, the uniform pressure loading rotates the connection structure only about one axis. For

a joint stiffness of $k_\beta = 100$ N-m/rad, the linear and non-linear solutions are shown in Fig. 5. The hinge rotations are identical in both $\beta = 56.499^\circ$ and 56.507° (exact solution) – the main difference is in the deformations of the brick elements. The nonlinear solution recovers a rigid body rotation about the flap hinge. The linear frequency, 1.239 Hz, is identical to exact solution $\sqrt{k_\beta/I_\beta}$. With a joint pitch spring of $k_\theta = 2.5$ N-m/rad, and pressure applied on only the top surface of leading edge bricks, the nonlinear solution is as shown in Fig. 6. Again, the joint rotation $\theta = 40.363^\circ$ and linear frequency 4.337 Hz are identical to exact solution. The nonlinear solution recovers a rigid body rotation about the torsion hinge.

Joint Modeling

For 3-D non-linear brick elements, joint modeling must ensure exact representation of a physical connection. This is dramatically different from the world of reduced order structural elements (beams and shells) where a connection is an idealization. This is because the physics of edge effects are non-existent in reduced order structural elements, whereas they are inherent in 3-D elements, and joint connections necessarily occur at the edges.

A joint can connect to any number of brick nodes on any number of bricks. All of these bricks must then be formulated as joint elements. The mathematical requirement is that a joint be connected to a minimum of 3 non-colinear nodes in order to transfer all rotations.




Figure 7: **Joint connected to an entire face of brick nodes occurring on a section; full face connection.**

The simplest joint model (baseline) is a full face connection as shown in Fig. 7. Here all elements on an en-




Figure 8: **Joint connected to a selected set of brick nodes representing a physical blade attachment; bolt attachment connection.**

tire face of the connection side structure are designated as joint elements and all nodes that lie on the face (9 nodes per element) as connection nodes. This model was used in the previous subsection. The formulation, however, is generic in that the joint elements can be embedded anywhere within a structure and any number of its nodes designated joint nodes. The joint elements need not be physically adjacent to a joint. Figure 8 shows an example of internal nodes of several bricks lying across the blade thickness connected to a joint. This represents a realistic bolt attachment type connection and will be studied in greater detail in the next subsection. A realistic torque tube pitch link connection is shown in Fig. 9. The joint here is connected to face nodes of several bricks that make up the torque tube. This model will be used later to analyze a bearingless rotor hub.

Under static loading, the rotations on the joint (linear or non-linear) are independent of the joint model. The end deflections and the internal stresses of the finite element structure that connects to the joint depend on the joint model. Under dynamic loading, these internal stresses govern the non-linear stiffness of the structure. This non-linearity associated with 3-D end stresses can have a dramatic impact in rotors because of the enormous centrifugal force field. This effect is described in the next subsection.

3-D Edge Effects in Rotors

Beam based rotor models with multibody joint connections at the root end neglect 3-D edge effects. This section demonstrates how this neglect can lead to significant discrepancy in torsion dynamics, and how 3-D models with realistic joint modeling can remedy this discrepancy.




Figure 9: Pitch link connection to torque tube for a bearingless rotor. Rigid connection from joint to connection points.

Mode No.	Hingeless	Articulated	Articulated in pitch
1	0.824	0.000	0.824
2	1.057	0.939	0.939
3	2.763	1.000	1.057
4	5.049	2.601	2.763
5	5.201	3.894	5.049
6	6.419	4.847	5.201
7	8.526	7.927	8.526
8	12.749	10.942	12.734

Table 1: Hingeless, articulated, and pitch-only articulated frequencies using beams. Torsion frequencies in bold.

Consider an uniform, untwisted rotor blade, of same dimensions and material properties that was verified earlier (see FEM formulation). Consider three boundary conditions, a hingeless, a fully articulated, and a pitch (only) articulated condition. The rotor frequencies using beams are given in Table 1. The pitch articulation, as expected, simply re-replaces the hingeless torsion frequencies with the articulated ones. The same frequencies calculated using bricks are given in Table 2. In the brick model, the hingeless and articulated boundary conditions are straight forward to implement, and are as shown in Fig. 10. The articulated condition is implemented as $u_1 = u_2 = u_3 = 0$ at a single node at the root boundary. This is referred to later as a boundary articulation. The hingeless and articulated frequencies are identical to beam frequencies as expected from the earlier verification (see Figs. 2 and 3). The pitch articulation, however, cannot be implemented in a straight forward manner, as there are no rotation variables in the bricks. The bricks must now be connected to a revolute joint first at the root end, and then articulation implemented on the joint vari-




Figure 10: Hingeless and articulated boundary conditions.




Figure 11: First torsion mode for an articulated rotor with zero pitch spring; 0.939/rev.

able. This is referred to later as a joint articulation. The rotor frequencies for pitch articulation shown in Table 2 uses a full face joint model (as in Fig. 7) to implement this connection. It is clear that this connection, to a joint with zero stiffness, does not recover the articulated torsion mode (with natural frequency of 0.939/rev) in bricks – unlike in beams where a multibody connection is bound to do so when interfaced with a joint with zero stiffness.

This artificial stiffening of torsion in 3-D is not an analysis error, but the result of a non-physical connection to bricks that are naturally equipped to capture the physics of 3-D end effects exactly. Unlike beam models, which idealize torsion as a separate variable, in 3-D, torsion arises naturally from a combination of inplane and transverse displacements. It will be shown that the inplane displacements at the edge of a rotor blade are substantial, and of an unique nature due to the action of centrifugal forces. By constraining every one of the




Figure 12: Internal axial deflections on hingeless and articulated rotors under pure rotation.

Mode No.	Hingeless	Articulated	Articulated in pitch
1	0.824	0.000	0.826
2	1.058	0.939	1.058
3	2.768	1.000	2.262
4	5.005	2.606	2.775
5	5.222	3.876	5.059
6	6.686	4.852	5.336
7	8.597	7.954	9.175
8	12.864	10.757	13.265

Table 2: Hingeless, articulated, and pitch-only articulated frequencies using multibody bricks. Torsion frequencies in bold.

end nodes to pure rotation, the full face joint model constrains these inplane 3-D end displacements, and in turn, torsion. The remedy is to use a joint model that represents a true physical connection, and therefore capture the physics of 3-D end effects precisely. In a true physical connection, the blade is bolted at the root end to an attachment piece and the attachment piece then connects to a bearing. The joint represents the bearing, and the joint model represents the connection provided by the attachment piece. Assuming the attachment piece is rigid (i.e. stiff), Fig. 8 is an illustration of such a connection. This is referred to as a bolt attachment model.

To study the 3-D edge effects, the same rotor is equipped with a grid that is now concentrated towards the root end. This does not affect the first modes, and the first three articulated frequencies remain 0, 0.939, and 1.0/rev – in lag, torsion, and flap. Figure 11 shows this torsion mode of interest along with the grid. The torsion frequency equals $\sqrt{(I_1 + I_2)/(I_2 - I_1)}$, same as the solution of beam torsion equation

$$(I_1 + I_2)\ddot{\phi} + \Omega^2(I_2 - I_1)\phi - (GJ\phi')' = 0$$

where I_2 and I_1 the chordwise and flapwise moments of




Figure 13: Internal inplane deflections of a hingeless rotor under pure rotation.




Figure 14: Internal inplane deflections of an articulated rotor; boundary articulation.

inertia and the root boundary condition is that of a free end, with $GJ\phi' = 0$.

For both the hingeless and articulated rotors, the internal transverse deflections (u_z) are negligible under pure rotation. The axial deflections are as in Fig. 12, with the span-wise variations at several chord stations plotted in the thickness-wise center plane. (The trends are similar in all thickness-wise planes.) The variation along mid-chord reaches zero at the root end, the other chordwise stations are symmetric about mid-chord. The variations are identical for the two rotors, with a constant offset, except at the root end of the articulated rotor. The axial deflections generate an inplane contraction via the Poisson's ratio effect, and these inplane deflections play a key role in 3-D edge effects. At the edge, these deflections are of a radically different nature between the hingeless and articulated rotors, as shown in Figs. 13 and 14 respectively. For over 95% of span, the variations are identical and exhibit a Poisson's contraction (positive deflection in trailing edge and negative deflection in leading edge) due to the axial elongation from centrifugal load-




Figure 15: **Root end close-up of internal inplane deflections of an articulated rotor; boundary articulation.**




Figure 17: **Internal inplane deflections of an articulated rotor; articulation via joint; joint model: full face connection.**




Figure 16: **Articulated rotor model using multi-body bricks; joint model: full face connection.**

ing. But at the inboard 5% near the root end, the free edge of the articulated rotor exhibits a dramatic local spike with cross-over characteristics. A close up view is shown in Fig. 15. Only at one node, as prescribed by the boundary condition, is the deflection zero.

Any blade attachment, regardless of its form, is bound to provide a constraint on these displacements, alter the internal stresses, and affect the non-linear torsion frequency. For example, the full face joint model, implemented as Fig. 16, produces an internal inplane deflection field as in Fig. 17, that bears little resemblance to the articulated rotor and appears closer to the hingeless boundary condition. The bolt attachment model, implemented as Fig. 18, applies the joint constraints only at locations where the blade bolts are expected to attach, and leaves the edge free. Therefore, the inplane deflec-




Figure 18: **Articulated rotor model using multi-body bricks; joint model: bolt attachment connection.**

Mode No.	Face interface	Bolt 20%-80%c	Bolt 30%-70%c	Bolt 37%-63%c
1	0.826	0.872	0.843	0.789
2	1.058	1.079	1.080	1.079
3	2.262	1.769	1.632	1.446
4	2.775	2.841	2.844	2.841
5	5.059	5.251	5.141	4.957
6	5.336	5.524	5.535	5.525
7	9.175	9.626	9.651	9.626
8	13.265	13.156	13.134	13.073

Table 3: Blade frequencies with revolute joint at root (pitch); zero joint stiffness; torsion frequencies in bold.




Figure 19: Internal inplane deflections of an articulated rotor; articulation via joint; joint model: bolt attachment connection; bolts at 2.5% R and 20%-80% c.

tions now retain the basic cross-over characteristics of the articulated rotor (Fig. 19), and consequently the resulting frequencies (Table 3) move closer to an articulated rotor. The table shows the results of bolt attachments at 2.5% R and at three sets of chordwise positions. The deflections for the three cases (boundary articulated, a full face joint, and bolt attachment joint at 2.5% R, 20%-80% c) are summarized for comparison in Fig. 20. It is clear that the different edge conditions produce identical internal deflections over the outboard 95% of the blade span. Yet, it is the inboard 5%, near the root end that has an important influence on the torsion frequency. As to be expected, the span-wise position of the attachment also matters. This variation is shown in Fig. 21. Clearly, when the attachment points lie too close to the edge to be physical, spurious frequencies can result (less than 0.5/rev).

In conclusion, we note that a beam based multibody model ignores the blade attachment entirely, and consequently neglects all of the 3-D edge effects. By connecting a joint directly to a beam and effectively implementing a boundary condition of the form $GJ\phi' = -k_\theta\phi$ on




Figure 20: Internal inplane deflections of an articulated rotor on leading and trailing edges compared using: (1) boundary articulation, (2) joint articulation with full face joint model, and (3) joint articulation with bolt attachment joint model.




Figure 21: First torsion frequency vs. spanwise location of bolt attachment points.

the torsion variable ϕ , it bypasses the dependence of ϕ on edge effects and ignores the non-linear edge stresses. From this study, it appears that no realistic connection can ever reproduce the ideal edge conditions assumed by beam based multibody models, and the error from this neglect is a significant one. Finally, it is the high centrifugal forcing that appears to give rise to the non-linear 3-D edge effects – a phenomena that is unique to rotors.

PARALLEL NEWTON-KRYLOV SOLVERS

The Parallel Newton-Krylov solvers are based on those originally reported in Refs. [2] and [14]. The main contribution here is the accommodation of multibody dynamics.

Each Newton iteration consists of a fully parallel lin-

ear solver based on iterative substructuring. In iterative substructuring, the substructure interiors are solved using direct factorization. This operation is naturally parallel. The substructure interfaces are solved iteratively, using Krylov updates, the building blocks of which are constructed using fully parallel substructure-by-substructure operations. The building blocks are: (1) residual calculation, (2) preconditioning of the residual, and (3) a matrix-vector multiplication procedure. The Krylov updates – Conjugate Gradient (CG) updates for symmetric systems and Generalized Minimum Residual (GMRES) updates for non-symmetric systems – are constructed using these building blocks.

The goal of iterative substructuring is to construct the building blocks in a scalable manner. This means if the substructures have an average size H , and the finite element mesh within each substructure has an average size h , then the condition number of the preconditioned interface problem must not grow with the number of substructures as long as the mesh within each substructure is refined to keep H/h constant. A large problem will then converge with the same number of Krylov updates (iteration counts) as a small problem. The preconditioner is then called an ‘optimal preconditioner’ and the solver is said to exhibit ‘optimal numerical scalability’.

The FETI-DP algorithm is such an iterative substructuring method. It can be constructed to guarantee optimal numerical scalability for problems governed by PDEs of up to 4th order and with heterogeneous properties.





Figure 22: Each figure is a top view of 4 neighboring substructures; at a dual node common to 4 substructures continuity can be enforced pairwise by (a) a minimum of 3 constraints across 3 pairs of substructures, (b) 4 constraints across 4 pairs and (c) a maximum of 6 constraints across 4 pairs.

The FETI-DP Algorithm

In the FETI-DP algorithm, the substructure interface is sub-divided into two categories: a selected set of corner nodes and a remaining set of non-corner nodes. The corner nodes are used to formulate a primal interface problem. Hence they are also termed primal nodes. The non-corner nodes are used to formulate a dual interface problem. Hence they are also termed dual nodes. In the primal interface problem, the variables (called primal



(a) 3 constraints

Figure 23: Each figure is a top view of 3 neighboring substructures; at a dual node common to 3 substructures continuity must be enforced pairwise by 3 constraints across 3 pairs of substructures.

variables) are the original finite element degrees of freedom. In the dual interface problem, the variables (called dual variables) are a set of auxiliary variables, that are not a direct subset of the original finite element degrees of freedom. Each dual variable is used to enforce continuity of the original finite element degrees of freedom across two substructures. The two interface problems are coupled, and the building blocks of the coupled dual-primal interface problem can be constructed in a fully parallel manner requiring communication only between the dual nodes of neighboring substructure — as long as the primal nodes are available in all. The primal problem is therefore solved in every processor and requires a global communication between all substructures. The primal nodes or corner nodes are the key to ensuring optimal numerical scalability. These form a coarse finite element representation of the problem, and ensure scalability by propagating local substructure information globally.

Each substructure interface node can be a face, edge, or a vertex node. A node that is common to two and only two substructures is a face node. A node that is common to at least three substructures is an edge node. Of these, those that occur at the end point of edges are vertex nodes. The edge and vertex nodes that are common to more than two substructures can be selected as corner nodes. This selection was used in our earlier work in Ref. [2]. This however leads to a large number of coarse nodes, and because the coarse problem require global communication, they limit the linear speed-up range for a given problem size. In this paper, only the vertex nodes are selected as corner nodes. This is a minimal selection as it excludes all edge nodes. An illustration is given later in the section on ‘partitioning and corner selection’.

To implement the minimal coarse problem, the FETI-DP solver must now treat all substructure edge nodes that connect to four substructures as dual nodes, as in Refs. [21, 22]. Each of these dual nodes must then be equipped with sufficient dual variables to enforce continuity of finite element degrees of freedom across, not two, but four substructures. As illustrated in Fig. 22, a minimum of three dual variables per nodal degree of freedom is required for this purpose, each enforcing con-

tinuity across a single pair of substructures. However, a maximum of six can be used leading to a set of multiple redundant dual variables. Unless otherwise mentioned, all subsequent results will use the full set of six dual variables.




Figure 24: **Partitioning of an articulated blade like structure containing an inboard multibody joint.**

FETI-DP with Joints

The joint DOFs are treated as internal nodes occurring within a substructure. The intent is to leave the substructure interfaces, and consequently the numerical scalability of iterative substructuring, unaffected by multibody dynamics. It also leaves unaffected the invertibility condition of the non-corner subdomain nodes as long as they are already selected to do so prior to inclusion of multibody dynamics. Thus, in the present formulation of the solver, grid generation and partitioning of the finite element structure remain unconstrained from the requirement of including multibody components.

The inclusion of multibody components as internal nodes leaves the FETI-DP algorithm unaffected but requires modification to the nodal ordering and connectivity of the partitioned subdomain nodes. The substructure nodes are still re-ordered as internals first (I^s), followed by the interface dual nodes (Γ_E^s), and then the interface primal nodes (Γ_C^s), where the subscript s denotes subdomain quantities, but now the internal nodes are re-arranged to order the multibody DOFs first followed by the rest. Placing the multibody DOFs ahead of the finite element nodes allows a simple shift in the existing nodal ordering and connectivity to be sufficient for including them within the solver. The joint connection nodes are placed at the end with boundary nodes (Γ_B^s).

Including multibody nodes as internal nodes, however, determines the rules based on which the original finite element structure is partitioned. For illustration, consider a simple blade like structure containing a joint




Figure 25: **Interface construction of a partitioned articulated blade like structure containing an inboard multibody joint.**

at an inboard section (Fig. 24). The presence of the joint prevents a straight forward 2-D partitioning of the entire structure without leaving the joint on a substructure interface. One alternative is a 1-D partition, but as shown earlier in Ref. [2], a 2-D partition is fundamentally superior in terms of efficiency. A simple solution is to carry out a 1-D partition locally near the joint, while performing a 2-D partition over the remainder of the blade. An example is shown in Fig. 24. The partitioning, however, dictates two changes to the construction of the solver. First, the subdomains that occur at the junction will contain dual nodes that must ensure continuity of displacements across an odd number of subdomains (3 in this case). Second, the primal node selection must be consistent across the two partitions. The first change implies that the subdomain that lies on the 1-D side of the junction will now contain special dual nodes that are to be equipped with multiple dual variables but are no longer geometric edges. The second change implies that it will now contain primal nodes that are not geometric vertices. Based on these rules, the interface construction is as shown in Fig. 25.

In order to compare scalability, the same beam, the same partition (as in Fig. 24), and the same algorithmic construction (as in Fig. 25) is used on two candidate structures – one in which a multibody joint is introduced at an inboard section and another in which it is not. The multibody joint model is chosen to be a full face model so as to introduce maximum discrepancy in the bandwidth and conditioning of the internal nodes between the two candidate structures. The convergence rate of the domain decomposition solver remains identical for both structures, as shown in Fig. 26. Both converge with nominally the same iteration count demonstrating that




Figure 26: Convergence of FETI-DP/CG solver with and without multibody dynamics; articulated rotor.




Figure 28: Partitioning and interface construction of a bearingless blade like structure containing an inboard multibody joint.




Figure 27: Illustration of a bearingless blade like structure containing an inboard multibody joint.

the condition number of the preconditioned interface has not been altered. This implies that if the original problem was scalable then the new problem with multibody dynamics remains scalable.

Next consider a simple prototype of a bearingless rotor, Fig. 27. The details of geometry and grid are given in the next section. A simple 1-D partitioning is carried out that ensures the pitch link joint is embedded entirely within the root substructure, see Fig. 28. Although 1-D, care is taken to avoid choosing a large number of corner nodes – a major drawback of a 1-D partitioning as presented in Ref. [2] – as only a few edge nodes are selected as corners (see subfigure in Fig. 28). The number of coarse nodes here are then exactly same as a 2-D partition. Again, two structures are considered, one in which the pitch link joint occurs as shown in the figure, and another without. Figure 29 shows that the rate of convergence is again similar, demonstrating the equivalence in the condition number of the preconditioned interface, and hence implying the same scalability.




Figure 29: Convergence of FETI-DP/CG solver with and without multibody dynamics; bearingless rotor.

In summary, including multibody dynamics within the domain decomposition solver sets important rules for partitioning the finite element structure. These rules in turn impact the algorithmic construction of the underlying parallel solver. However, once the rules are followed, the present method of including multibody dynamics leaves the subdomain interfaces, and consequently the scalability, of the domain decomposition solver fundamentally unchanged.

Parallel CG and GMRES Updates

In addition to the communication required by FETI-DP in constructing the building blocks, the CG and GMRES updates require additional processor synchronization points of their own. These must be minimized to prevent high communication costs diminishing scalability of the parallel implementation regardless of the numerical scalability of the underlying algorithm.

A Conjugate Gradient (CG) update requires three processor synchronization points – vector inner products that require global communication including a norm calculation to determine the stopping criteria. The total number can be reduced to just one using advanced norm estimation techniques [23, 24]. This has not been included at present. The requirement is more severe for the GMRES update and is more relevant to rotary wing structures due to its non-symmetric nature.

A Generalized Minimum Residual (GMRES) update incurs significantly more communication cost than a CG update. At the heart of a GMRES update is the Arnoldi algorithm. To solve $Ax = b$, it constructs m orthonormal basis vectors $V_m = [v_1, v_2, \dots, v_m]$ spanning the m -dimensional Krylov subspace $K_m(A, r_0) = \text{span}(r_0, Ar_0, \dots, A^{m-1}r_0)$, where $r_0 = b - Ax_0$ and x_0 is the current estimate of the solution, and a matrix \bar{H}_m of size $(m+1) \times m$ the top $m \times m$ block of which is an upper Hessenberg matrix H_m . The construction of each vector requires orthogonalization with respect to every one of the previous. Traditionally, a Modified Gram-Schmidt procedure is preferred for this orthogonalization step because of its numerical stability over Classical Gram-Schmidt. However it requires as many as m synchronization points compared to only one in Classical Gram-Schmidt. In this study we implement a Reorthogonalized Classical Gram-Schmidt procedure that produces orthogonalization superior to Modified Gram-Schmidt while requiring only two synchronization points [25, 26].

COMPONENTS OF MULTIBODY-FEM ANALYSIS




Figure 30: **Planform of a prototype rotor blade used in this study; $c = 0.53$ m.**

3-D Geometry and Grids

Geometry and grids are critical components of a 3-D rotor analysis, but are not the present focus of this work. It is assumed that suitable geometry and grid generators will be available to the solver from other sources. Yet, for the purposes of solver development, and for understanding the solver requirements for a multibody based 3-D FEM analysis, a simple grid generator, a component merger, and a joint locator are developed.

The grid generator can discretize only one continuous structure at a time and assumes that the cross-sectional discretization remain same along span. Within




Figure 31: **Components of a hingeless-bearingless rotor blade.**

this assumption, it is easy to accommodate arbitrary airfoil shapes, twist, planform, and advanced geometry tips. The structure can be solid or shell-like with several layers of bricks. Different element groups can be prescribed different material properties (e.g. spar, skin, and honeycomb). The component merger can merge two or more gridded structures but assumes the same grid resolution at the connection interface. The merger conforms geometry as well as nodal connectivities at the interface. The joint locator identifies the multibody joint connection nodes and adds them to boundary nodes for elimination.

Grid	$n_1 \times n_2 \times n_3$	DOFs
Small scale		
Hingeless 1	$96 \times 4 \times 2$	25,920
Hingeless 2	$48 \times 4 \times 4$	25,920
Hingeless 3	$64 \times 4 \times 4$	34,560
Large scale		
Hingeless	$128 \times 12 \times 12$	480,000
Articulated	$128 \times 12 \times 12$	480,000
Bearingless	$127 \times 12 \times 12 + 136$	482,271

Table 4: **3-D FEM hingeless rotor grids for scalability study.**

Three prototype rotor blades are considered: (1) a hingeless blade, (2) an articulated blade with coincident flap-lag-torsion hinge, and (3) a bearingless blade. The hingeless blade is a single component structure. The articulated and bearingless blades are multibody structures.

The nominal blade geometry is shown in Fig. 30. It contains a generic symmetric airfoil of 5% thickness at every radial station. The planform is generic with a




Figure 32: A hingeless rotor blade prototype with $128 \times 12 \times 12$ elements (every 3 span stations shown); 0.48 M FEM degrees of freedom




Figure 33: Cross-section of prototype blade showing 12×12 bricks with 625 nodes; exaggerated vertical scale.

sweep of 20° outboard from 95% span station. The articulated blade has an identical geometry, except for a spherical joint located at 5% R to provide articulation. The bearingless blade contains multiple flexible components at the root end as illustrated in Fig. 31 along with a pitch link connection to the torque tube on the retreating side. The joint allows vertical, inplane, and pitching motion while constraining all others.

Each finite element can accommodate its own material model and ply direction but for the purposes of scalability and timing simple isotropic properties suffice: $E = 73$ GPa; $\nu = 0.3$; and $\rho = 2700$ kg/m³. The rotational speed is a steady $\Omega = 27$ rad/s. With $c = 0.53$ m, these values generate typical stiffness and inertia of soft in-plane rotors for the hingeless blade. No attempt is made to place the sectional offsets at quarter-chord for any of the configurations.

The simple geometry of the hingeless blade is ideally suited for scalability study as it can be partitioned uniformly into many substructures containing the same number of elements. The numerical scalability of the solver can then be examined without the practical constraints of load balancing. We consider three small scale problems for the scalability study and three large scale problems for timing study as listed in Table 4. n_1 , n_2 ,




Figure 34: An articulated rotor blade prototype with $128 \times 12 \times 12$ elements (every 3 span stations shown); 0.48 M FEM degrees of freedom, with coincident flap, lag, torsion articulation at 5% R

and n_3 are numbers of elements along span, chord, and thickness. By small scale, we mean sizes that can also be analyzed on a single processor for a consistent comparison of parallel speed-up.

For practical applications, not just scalability, but the actual run times are of prime importance. The large scale problems are designed for this purpose. The increase in problem size is achieved primarily by increase in cross sectional resolution. The largest problem size for the hingeless rotor consists of 0.48 million (M) DOFs. For this size, the discretized blade and the cross section are shown in Figs. 32 and 33 respectively.

The articulated and bearingless rotors considered are both large scale problems. The presence of multi-body components in both calls for special partitioning and load balancing. The articulated blade contains the same number of DOFs as the hingeless blade with the addition of 6 joint DOFs. The bearingless rotor contains 482,271 DOFs out of which 1,326 are in the torque tube including the 6 joint DOFs and 585 DOFs are in the flex-element. The discretized configurations are shown in Figs. 34 and 35.

Partitioning and Corner Selection

The inclusion of multibody dynamics require that partitioning be carried out satisfying two criteria. First, that the multibody components are entirely embedded inside substructures, i.e., both the joint attachment node as well as all the joint connection nodes are part of bricks that are contained within the same substructure. Second, that the total DOFs are still evenly distributed across substructures. Once these criteria are met, the partitioner then performs the three tasks: (1) re-orders sub-




Figure 35: A hingeless-bearingless rotor blade prototype; a total of 0.48 M FEM degrees of freedom for entire structure, with pitch link articulation.




Figure 37: Blade partitioned into a combination of 1-D and 2-D partitions; $2 \times 1 + 7 \times 2 = 16$ substructures shown for illustration.




Figure 36: 3-D FEM of a hingeless rotor blade using isoparametric brick elements; blade partitioned into 8×2 substructures for illustration.




Figure 38: A bearingless blade partitioned into 16×1 substructures shown for illustration.




Figure 39: A typical substructure showing baseline coarse problem selection; circles are dual interface nodes, squares are primal coarse nodes.




Figure 40: A typical substructure showing minimal coarse problem selection; circles are dual interface nodes, squares are primal coarse nodes.

structure nodes and element connectivity, (2) selects corner nodes, and (3) constructs substructure to substructure communication maps.

The node re-ordering brings the interior nodes first, followed by interface nodes, and then the boundary nodes. Among interior nodes, the multibody DOFs are arranged first, followed by the finite element nodes. The boundary nodes are augmented with the joint connection nodes. The interface nodes consist of face, edge, and vertex nodes. These are then separated into corner and non-corner nodes for treatment as primal and dual interface nodes respectively. All dual interface nodes that occur at the edges are identified as special and equipped with multiple dual variables as required.

For the hingeless blade, a typical partitioning is illustrated in Figure 36. Selection of corner nodes is the most important requirement and must be performed in an intelligent manner. First, the selection must ensure null kernels in every substructure, i.e. constrain rigid body motion by ensuring that the non-corner restriction of the stiffness matrix is invertible. Second, it must be as small as possible, enough just to provide global error propagation but no larger. A selection containing all of the edge and vertex nodes common to more than two substructures (see Fig. 39) was used in our previous study [2] and is referred to in this paper as the baseline coarse prob-




Figure 41: Substructure merging 1-D and 2-D partitioned domains showing coarse problem selection; circles are dual interface nodes, squares are primal coarse nodes.




Figure 42: A typical substructure of a bearingless rotor; circles are dual interface nodes, squares are primal coarse nodes.

lem. The selection studied in this paper contains only a subset of these corner nodes, and consists only of the vertex nodes that lie at the end of the edges (see Fig. 40). This is referred to as the minimal coarse problem. Its size is now independent of the cross sectional grid and is at the most 8 per substructure. Note that the vertices that occur at the boundaries of the structure must also be included, even though they are common to only two substructures, to satisfy the first criteria of null kernels. Otherwise, the substructures at the tip end will contain rigid body rotational modes making them non-invertible.

For the articulated blade, a typical partitioning is a combination of 1-D and 2-D partitions as illustrated in Fig. 37. The combined 1-D and 2-D partitioning for the articulated blades generates similar substructures on the 2-D side. The substructures on the 1-D side, however,

require a different selection of the coarse problem. This selection is determined by the coarse problem of the substructure at the junction. Here, the coarse problem must maintain consistency with the 2-D partitions, as shown in Fig. 41. The coarse problem on the 1-D partitions then follows from the junction substructure.

For the bearingless blade, a typical partitioning is a 1-D partition as shown in Fig. 38. The 1-D partitioning generates typical subdomains as shown in Fig. 42. Note that the coarse selection here is superior to the 1-D partitioning originally shown in Ref. [2]. The coarse problem here is selected in the same manner as that of the 1-D partitions on the inboard sections of the articulated rotor.

The substructure to substructure connectivity needs to be calculated only once. Each substructure creates a destination and a reception map. The former contains the substructures to which quantities are to be sent, and the corresponding destination node numbers. The latter contains the substructures from which quantities are to be received, and the corresponding recipient node numbers. The dual nodes that lie on the edges communicate with four neighboring substructures. The dual nodes that lie on the faces communicate only with two neighboring substructures.

Hover and Forward Flight Prototypes

The hover prototype simply solves for steady blade response at a fixed collective of 10° with pressure airloads of 100 N/m^2 (418 lb/ft radial distribution) on the top surface. The airloads have the non-linear characteristics of a follower force. The non-linear solution procedure uses Newton-Raphson outer iterations. Within each iteration, the implicit FETI-DP inner solver uses CG updates. A CG update is adequate in ideal hover as the stiffness matrix is symmetric. The initial iterations converge the structural non-linearities associated with rotation. Once converged, the airloads are imposed. The virtual work during each airload iteration is calculated based on the previous iteration deformation state.

The transient forward flight prototype uses a Newmark scheme with a 5° azimuth step. The dynamic stiffness is now non-symmetric, therefore, the inner Krylov solver uses a GMRES update. For purposes of scalability study, the response for a single time step suffices, as the structure of the dynamic stiffness matrix remains same for all. We consider the following dimensions of Krylov subspace: $m = 30, 40$, and 50 , deemed more than adequate for large scale problems. Note that increasing m improves efficiency (faster convergence) at the cost of reduced scalability (greater communication).

SCALABILITY AND TIMING OF 3-D ROTOR ANALYSIS

The scalability study is carried out on the hingeless rotor with small scale problem sizes. The hingeless rotor can be partitioned into a large number of substructures without load imbalance. The small scale problems can also be analyzed on a single processor (without memory overflow) allowing the calculation of parallel speed-up. Once scalability is established, the timing study is then carried out for all of the three rotors of large scale problem size. Note that it has already been shown that the scalability of the solver remains same with or without multibody dynamics. The main effect of multibody dynamics is to set new rules for partitioning within which load balancing must be carried out carefully. For the articulated and bearingless rotors, the substructures are specially partitioned to confirm to these requirements.

Scalability Study

First, the study is conducted on a local unix cluster of 2.2 GHz dual core AMD Opteron processors. This to compare present results consistently with those reported earlier in Ref. [2]. Subsequently all computations are carried out on an Army DoD Supercomputing Resource Center (DSRC) cluster of 3.0 GHz dual core Intel Woodcrest processors. All times are wall clock times.

Consider the hingeless rotor of size $48 \times 4 \times 4$, partitioned into $n_s = 8 \times 2 = 16$ substructures (as in Fig. 36). The FETI-DP/CG (single Newton iteration in hover) solver times on a single processor for the baseline and minimal coarse problem implementations are compared in Fig. 43. It is clear that the optimal number of substructures — number of substructures for which the solver time is minimum — is extended by the minimal coarse problem. For the problem of size $48 \times 4 \times 4$ the baseline coarse node selection (as in Fig. 39) produces an optimality at 24 substructures whereas the minimal coarse node selection (as in Fig. 40) produces an optimality at 48 or more substructures. Similarly, for the problem of size $96 \times 4 \times 2$, the optimality is extended from 32 to 64 substructures.

n_s	FE	Sub. LU	Coarse problem	FETI	Solver total
8	198	453	125	517	1099
12	197	257	101	398	758
16	193	174	99	334	611
24	191	101	149	258	510
32	190	67	279	222	569
48	190	35	866	186	1088

Table 5: Solver time (s) vs. number of substructures n_s with baseline coarse problem; single processor; $48 \times 4 \times 4$ elements.

The reason behind this extension is clear from the




Figure 43: Solver time (s) vs. number of substructures for calculations on a single processor; $48 \times 4 \times 4$ elements; hover.




Figure 45: Parallel speed-up for calculations on multiple processors; each substructure on each processor; $48 \times 4 \times 4$ elements; hover.




Figure 44: Solver time (s) vs. number of substructures for calculations on a single processor; $96 \times 4 \times 2$ elements; hover.




Figure 46: Parallel speed-up for calculations on multiple processors; each substructure on each processor; $96 \times 4 \times 4$ elements; hover.




Figure 47: Solver time (s) vs. number of substructures for calculations on a single processor; three problem sizes; hover.




Figure 49: Solver time (s) vs. number of substructures for calculations on a single processor; three problem sizes; forward flight.




Figure 48: Parallel speed-up for calculations on multiple processors; three problem sizes; each substructure on each processor; hover.




Figure 50: Parallel speed-up for calculations on multiple processors; three problem sizes; each substructure on each processor; forward flight.

n_s	FE	Sub. LU	Coarse problem	FETI	Solver total
8	198	496	32	601	1134
12	198	290	23	479	796
16	193	204	19	438	664
24	192	124	15	346	487
32	191	86	14	297	400
48	191	51	20	260	333
96	190	20	94	546	662

Table 6: **Solver time (s) vs. number of substructures n_s with minimal coarse problem; single processor; $48 \times 4 \times 4$ elements.**

detailed break-up of solver timings for the baseline and the minimal coarse problem implementations and are given in Tables 5 and 6. In the tables, ‘FE’ refers to the time taken to construct the structural matrices. ‘Solver total’ refers to the total solver time. The two together constitute the total simulation time. ‘Solver total’ consists of three parts: (1) ‘Substructure LU’ time, which refers to the substructure factorization, (2) ‘Coarse problem’ time, which refers to the coarse problem factorization, and (3) the ‘FETI-DP’ time, refers to the Krylov solver time including residual, preconditioner, and matrix-vector multiplies. The tables show that the dramatic reduction in coarse problem time and the delay in its growth leads to a significantly higher substructure optimality for the same problem size. This has important ramifications for scalability and timings for the parallel implementation.

The parallel implementation solves each substructure on a separate processor. To calculate parallel speed-up, the parallel solver time is compared with the serial solver time with the same number of substructures as the parallel solver. This ensures that computations of the same complexity are compared and that the speed-up is not contaminated with the benefits of substructuring itself.

The parallel speed-up for the two problems are shown in Figs. 45 and 46. In each figure, the speed-up obtained from the two coarse node selections are compared. It is clear that the minimal coarse node selection extends the linear speed-up range to a greater number of processors. Thus, for a given problem size, the minimal selection enables the fastest parallel solver time. From Fig. 45, the problem of size $48 \times 4 \times 4$ that could be solved in 21s using 24 processors, but no faster, can now be solved in 7s using 48 processors. The detailed break-up of the parallel solver times is given in Table 7.

Similarly, from Fig. 46, the problem of size $96 \times 4 \times 2$ that could be solved in 11s is now solved in 6s. However, for this problem the optimality is not yet reached with the available 48 processors. In order to study the full scalability range, all calculations are re-performed on the DSRC cluster, where more processors are available.

n_p	FE	Sub. LU	Coarse problem	FETI	Solver total
8	24	66	4.18	67	137
12	16	26	1.97	34	62
16	12	13	1.19	21	35
24	8.2	5.5	0.68	11	18
32	6.1	2.9	0.54	7.6	11
48	4.2	1.2	0.69	5.2	7

Table 7: **Solver time (s) vs. number of processors n_p with minimal coarse problem; $48 \times 4 \times 4$ elements.**

Henceforth, all studies are conducted on this platform. Figures 47 and 48 show the single processor timings and parallel speed-up respectively of the same problems. An additional problem of size $64 \times 4 \times 4$ elements is considered which could be partitioned into 128 substructures and analyzed on 128 processors. Even though the actual timings are significantly superior on this platform (5–10 times faster), the conclusions on scalability remain the same. The two problems of sizes $96 \times 4 \times 2$ and $64 \times 4 \times 4$ elements that have optimality of 64 show linear speed-up up to 64 processors, the problem of size $48 \times 4 \times 4$ that has optimality of 48 shows linear speed-up up to 48 processors. The solver times for serial and parallel computations for the problem of size $64 \times 4 \times 4$ are documented in Tables 8 and 9 respectively.

n_p	FE	Sub. LU	Coarse problem	FETI	Solver total
8	30	79	9.7	298	388
16	24	30	5.2	199	234
32	23	12	3.0	139	154
64	21	5.3	6.8	124	136
128	21	2.7	65.5	349	418

Table 8: **Solver time (s) vs. number of substructures n_s with minimal coarse problem; $64 \times 4 \times 4$ elements.**

n_p	FE	Sub. LU	Coarse problem	FETI	Solver total
8	2.5	11.95	1.63	40	53.4
16	1.1	1.59	0.48	12	14.4
32	0.55	0.33	0.16	4	4.52
64	0.27	0.08	0.16	1.8	2.04
128	0.16	0.02	0.78	3.4	4.23

Table 9: **Solver time (s) vs. number of processors n_p with minimal coarse problem; $64 \times 4 \times 4$ elements.**

The conclusions drawn on substructure optimality and parallel speed-up using the FETI-DP/CG solver is carried over to the FETI-DP/GMRES solver. Figures 49 and 50 show the single processor timings (single Newton

iteration of a single time step in forward flight) and parallel speed-up respectively. For these results, the GMRES solver uses a restart parameter of $m = 30$, and a Classical Gram-Schmidt with Re-orthogonalization based Arnoldi algorithm (see Ref. [2]). The actual timings are lower because the convergence criteria is set to 10^{-8} , as compared to 10^{-12} for the CG, due to the oscillatory nature of residual convergence beyond this value.

Timing Study for Large Scale Problems

It was demonstrated in the last section that the implicit parallel solvers developed using the FETI-DP method of iterative substructuring can solve hover and forward flight response in a scalable manner. For example, each Newton iteration of a 34,560 DOFs problem could be solved 64 times faster on 64 processors than on a single processor. For the solution of large scale problems, not just scalability but actual solver timings are of equal importance. By extending substructure optimality and linear speed-up to as high a processor number as possible, the minimal coarse problem now enables the benchmarking of actual solver timings on the three large scale rotor prototypes.

Each of the prototypes contain around 0.48 M DOFs. The main blade contains a cross-sectional resolution of 12×12 second order elements with a total of 25×25 nodes. The FETI convergence criteria for all cases are set to 10^{-6} for the preconditioned residual – more than adequate as the absolute residuals are always lower than this value.

For the hingeless rotor, the blade is discretized into $64 \times 2 = 128$ substructures and analyzed on 128 processors. Even though the substructure optimality of this model is expected to be far greater than this number, the partitioner is limited at present to spanwise and chordwise partitions, with no partitioning across thickness. At this level of decomposition, each substructure contains two layers of bricks each. The solver times for a single Newton iteration is shown in Table 10. The FETI-DP/CG solver is used on the symmetric stiffness matrix corresponding to hover. The FETI-DP/GMRES solver is used on the non-symmetric stiffness matrix corresponding to a single time step of implicit Newmark for transient forward flight. The forward flight cases converge faster because the mass matrix improves the condition number of the dynamic stiffness matrix leading to lesser number of iterations. The iteration count can be reduced further by using a greater value of restart parameter m . The consequent increase in communication, however, does not appear to incur a penalty as the solver time follows the same trend as iteration count. Henceforth $m = 40$ is used as baseline.

The number of dual variables per edge corner (n_λ) has an important effect on solver time. Four variables per edge corner ($n_\lambda = 4$) is considered baseline in this study. The variation from a minimum of 3 to a maxi-

Solver type	FE	Sub. LU	Coarse problem	FETI	Solver total	Iter.
CG	2.9	35.4	3.14	220	258	509
GM30	2.9	35.9	3.15	142	180	325
GM40	2.9	35.5	3.14	135	173	309
GM50	2.9	35.4	3.14	130	168	296

Table 10: **Solver times (s) for FETI-DP/CG and FETI-DP/GMRES ($m = 30, 40, 50$) prototypes; analysis of 0.48 M hingeless model on 128 processors, each substructure on each processor.**

imum of 6 is shown in Table 11. In general, increase in number of dual variables leads to faster convergence but at a greater communication cost. From Table 11 however communication cost is not a concern — iteration count and solver times both show the same trends. It is clear that more than 3 is desired and 4 is close to optimal – hence chosen as baseline. 5 is not preferred as one out of the two cross directions (see Fig. 22) must be picked arbitrarily.

Dual variable per edge corner	GM30	GM40	GM50
3	252 (489)	225 (428)	220 (416)
4	180 (325)	173 (309)	167 (296)
5	183 (327)	159 (274)	164 (285)
6	202 (366)	183 (327)	179 (314)

Table 11: **Solver times (s) and iteration count (in brackets) vs. number of dual variables per edge corner; FETI-DP/GMRES with $m = 30, 40, 50$, analysis of 0.48 M model on 128 processors, each substructure on each processor.**

For the bearingless rotor, the structure is partitioned into a total of 1×64 substructures. The first substructure from the root end contains a total of $2 \times 4 \times 6 = 48$ bricks on the flex-element, 88 bricks on the torque tube, and 144 bricks on the main blade – a total of 276 bricks. The remaining substructures contain 288 bricks each. This is the closest load balancing that can be achieved with the present partitioner. That the same problem size can only be partitioned now into half the number of substructures is due to the presence of multibody dynamics. However, the problem stems not from multibody dynamics, but from the lack of effective partitioning at present. An effective partitioner is one that will decompose each of the root components separately. The only requirement then would be that the multibody components be placed entirely within one of the torque tube substructures. The solver times for three loading cases are shown in Table 12. Apart from hover and forward flight a non-rotating static condition is also shown (with pressure loading reduced to 0.1 N/m^2). Iterative solvers have problem dependent convergence. In general, the non-rotating case has the fastest convergence. Introduc-

ing rotation introduces stiffness which increases convergence time. The dynamic stiffness with inertial terms in general converges faster for reasons given above.

Analysis type	FE	Sub. LU	Coarse problem	FETI	Solver total
Static loading	5.6	202	6.7	102	310
Hover	5.7	204	7.3	302	513
Forward flight	5.7	204	7.3	235	447

Table 12: **Solver times (s) for FETI-DP/CG and FETI-DP/GMRES ($m = 40$) prototypes; analysis of 0.48 M bearingless model on 64 processors, each substructure on each processor.**

For the articulated rotor, the structure is partitioned into a total of $(1 \times 4) + (2 \times 30) = 64$ substructures. The first 4 from the root end are part of 1-D partitioning. Each contain $2 \times 12 \times 12 = 288$ bricks. The fourth substructure carries the multibody joint (at 5% R). It is also the connecting substructure to the rest of the blade which is partitioned using 2-D partitioning. Each substructure here contains $4 \times 12 \times 6 = 288$ bricks. Thus, the substructures are all load balanced – except for the fourth which contains 3 additional DOFs, a difference of no consequence. The solver times for three loading cases are shown in Table 13. For the articulated rotor, the convergence times show a different trend. All three loading conditions show similar time for convergence, and in general converges faster compared to the hingeless and bearingless rotors. This is because of the hinge. The stiffness of the structure is now determined primarily by the hinge articulation that makes it less stiffer than the other configurations.

Analysis type	FE	Sub. LU	Coarse problem	FETI	Solver total
Static loading	5.6	189	6.4	116	312
Hover	5.6	188	6.3	115	310
Forward flight	5.6	186	6.2	125	317

Table 13: **Solver times (s) for FETI-DP/CG and FETI-DP/GMRES ($m = 40$) prototypes; analysis of 0.48 M articulated model on 64 processors, each substructure on each processor.**

The variation of solver time with number of processors cannot be compared fairly across different problem types, but it is clear that use of 64 processors over 128, approximately doubles the solution time. The comparison is somewhat meaningful between the hingeless and bearingless rotors as their boundary conditions are more similar to each other as compared to the articulated rotor. The bearingless rotor on 64 processors require 513 s (hover) and 447 s (forward flight) compared to 258 s (hover) and 173 s (forward flight) for the hingeless rotor on twice the number of processors — a factor of 2 and 2.6 respectively.

CONCLUDING OBSERVATIONS

The main objective of this paper was to provide a unified formulation for including multibody dynamics within a 3-D FEM based nonlinear finite element analysis for rotors and to devise and demonstrate a parallel solution procedure that accommodates multibody dynamics while maintaining scalability. To this end, special multibody brick elements were formulated in this study. The fundamental importance of precise joint modeling in connection with non-linear 3-D finite elements was highlighted. The physics of non-linear edge effects that require 3-D modeling was investigated. Then, a method to accommodate multibody degrees of freedom within fully parallel Newton Krylov iterative substructuring solvers was devised. It was demonstrated that the method leaves the numerical scalability of the original algorithm un-affected. Finally, three large scale prototype rotor configurations – one hingeless, one fully articulated, and one bearingless, containing up to 0.48 M DOFs and multibody hub components – were studied using the 3-D FEM multibody analysis, on up to 128 processors, for both hover and forward flight type calculations. Based on this study, the following key conclusions are drawn.

1. An unified 3-D FEM multibody dynamic analysis for rotor can indeed be carried out in a fully parallel and scalable manner – both in hover and forward flight.
2. An iterative substructuring algorithm that is scalable for pure finite element structural analysis can easily be tailored to accommodate multibody dynamics and still retain its scalability, as long as the multibody components are embedded wholly within substructures, leaving the interfaces free.
3. The above requirement, however, sets secondary rules on partitioning that calls for changes to the construction of the solver. In addition, it calls for a good partitioner for efficiently solving the problem on a large number of processors.
4. 3-D brick elements in connection with multibody modeling open opportunity to predictions of non-linear 3-D edge effects — so far unmodeled by current generation beam based dynamic analysis. This is a non-linear phenomena that is unique to rotors and appears to have a strong influence on torsion dynamics. It also highlights the need for precise modeling of joints to avoid spurious results.
5. Realistic rotor geometries containing up to half a million 3-D FEM multibody degrees of freedom can be solved in around 250 s – 500 s on 128 – 64 processors for a single Newton iteration of prescribed aerodynamic forcing.

Outlook and Future Research

The above timings (see last conclusion) translate approximately into 1 hr for a hover solution (with 20 Newton iterations for tight convergence, see Ref [2]) and 20 minutes per time step (5°) in forward flight (assuming 5 Newton iterations). It is clear that high fidelity 3-D structures cannot be enabled or driven without High Performance Computing, and without 3-D geometry and grid tools. But it is also clear, based on the research documented here and its companion paper earlier [2], that it is technologically within reach. Today, isolated rotor RANS CFD calculations, alone, require 20–40 hrs on 128–64 processors for hover computations, and about 40 mins per 5° azimuth (assuming 0.25° time steps for stability) in forward flight. Therefore, given the resources, there is little reason not to use 3-D structures for high fidelity analysis. But structures is fundamentally different from fluids — in physics, in mathematics, in numerics, and in the solutions sought. Mindful of these, we suggest a list of future directions of research subdivided into three categories.

1. *Fundamental*: Periodic dynamics – scalable domain decomposition in space-time for direct extraction of periodic dynamics.
2. *Applied*: Dual node based coarse problem augmentation with 3-D partitioning. Large-scale eigensolution; 3-D to 3-D fluid-structure interface. Reduced order aerodynamics and trim solution.
3. *Applications*: Integration of 3-D solid geometry and grid tools; Analysis of production rotors with multibody hub components using design geometries and properties, beginning with hover.

ACKNOWLEDGMENTS

This research is conducted at the U. S. Army Aeroflightdynamics Directorate as part of the HPC Institute of Advanced Rotorcraft Modeling and Simulation (HI-ARMS) program supported by the U. S. DoD HPC Modernization Program Office, led by Dr. Roger Strawn. The authors acknowledge his support for this research. The authors thank Prof. Charbel Farhat, Stanford University, for his insights on domain decomposition.

REFERENCES

- [1] Johnson, W. and Datta, A., “Requirements for Next Generation Comprehensive Analysis of Rotorcraft,” American Helicopter Society Specialist’s Conference on Aeromechanics, San Francisco, CA. Jan. 23-25, 2008.
- [2] Datta, A. and Johnson, W., “Three-dimensional Finite Element Formulation and Scalable Domain Decomposition for High Fidelity Rotor Dynamic
- Analysis,” American Helicopter Society Annual Forum 65, Grapevine, Texas, May 27–29, 2009.
- [3] Johnson, W., “Rotorcraft Dynamics Models for a Comprehensive Analysis,” American Helicopter Society Annual Forum 54, Alexandria, VA, May 20–22, 1998.
- [4] Bauchau, O. A., Bottasso, C., and Nikishkov, Y., “Modeling Rotorcraft Dynamics with Finite Element Multibody Procedures,” *Mathematical and Computer Modeling*, Vol. 33, (10-11), May 2001, pp. 1113-1137.
- [5] Saberi, H., Khoshlahjeh, M., Ormiston, R. A., and Rutkowski, M. J., “Overview of RCAS and Application to Advanced Rotorcraft Problems,” American Helicopter Society Aeromechanics Specialists Conference, San Francisco, CA, January 2004.
- [6] Shabana, A. A., Bauchau, O. A., and Hulbert, G. M., “Integration of Large Deformation Finite Element and Multibody System Algorithms,” *Journal of Computational and Nonlinear Dynamics*, Vol. 2, (4), October 2007, pp. 351–359.
- [7] Gerstmayr, J. and Schoberl, J., “A 3D Finite Element Method For Flexible Multibody Systems,” *Multibody System Dynamics*, Vol. 14, (4), May 2006, pp. 305–320.
- [8] Ibrahimbegovic, A., Taylor, R. L., and Lim, H., “Non-linear Dynamics of Flexible Multibody Systems,” *Computers and Structures*, Vol. 81, (12), May 2003, pp. 1113–1132.
- [9] Orden, J. C. G. and Goicolea, J. M., “Conserving Properties in Constrained Dynamics of Flexible Multibody Systems,” *Multibody System Dynamics*, Vol. 4, (2-3), August 2000, pp. 225–244.
- [10] Berdichevsky, V. L., “Variational-Asymptotic Method of Constructing a Theory of Shells,” *Journal of Applied Mathematics and Mechanics*, translated from *Prikladnaya Matematika i Mekhanika*, Vol. 43, (4), 1979, pp. 664–687.
- [11] Hodges, D. H., *Nonlinear Composite Beam Theory*, AIAA, Reston, VA, 2006.
- [12] Datta, A., Nixon, M. and Chopra, I., “Review of Rotor Loads Prediction with the Emergence of Rotorcraft CFD,” *Journal of the American Helicopter Society*, Vol. 52, (4), October 2007, pp. 287–217.
- [13] Yamauchi, G. K. and Young, L. A., Editors, “A Status of NASA Rotorcraft Research,” NASA/TP-2009-215369, 2009.
- [14] Datta, A. and Johnson, W., “Large-Scale Domain Decomposition For A Scalable, Three-dimensional

- Brick Finite Element Based Rotor Dynamic Analysis,” American Helicopter Society Specialists’ Conference on Aeromechanics, San Francisco, CA, January 20-22, 2010.
- [15] Farhat, C., Lesoinne, M., and Pierson, K., “ Scalable Dual-Primal Domain Decomposition Method,” *Numerical Linear Algebra with Applications*, Vol. 7, (7-8), 2000, pp. 687–714.
 - [16] Farhat, C., Lesoinne, M., LeTallec, P., Pierson, K., and Rixen, D., “FETI-DP: A Dual-Primal Unified FETI Method - Part I: A Faster Alternative to the Two-level FETI Method,” *International Journal of Numerical Methods in Engineering*, Vol. 50, (7), pp. 1523–1544, 2001.
 - [17] Klawonn, A., and Rheinbach, O., “Robust FETI-DP Methods for Heterogenous Three Dimensional Elasticity Problems,” *Computer Methods in Applied Mechanics and Engineering*, Vol. 196, (8), 2007, pp. 1400–1414.
 - [18] Bathe, K., *Finite Element Procedures in Engineering Analysis*, Prentice-Hall, Inc., NJ, 1982.
 - [19] Zienkiewicz, O. C., and Taylor, R. L., *The Finite Element Method for Solid and Structural Mechanics*, Elsevier Butterworth-Heinemann, Oxford, UK, Sixth edition, 2006.
 - [20] Zienkiewicz, O. C., *The Finite Element Method in Structural and Continuum Mechanics*, McGraw-Hill, London, U.K., 1967.
 - [21] Rixen, D. J. and Farhat, C. “A Simple and Efficient Extension of a Class of Substructure Based Preconditioners to Heterogenous Structural Mechanics Problems,” *International Journal for Numerical Methods in Engineering*, Vol. 44, (4), pp. 489–516, January 1999.
 - [22] Klawonn, A. and Widlund, O., “FETI and Neumann-Neumann Iterative Substructuring Methods: Connections and New Results,” *Communications on Pure and Applied Mathematics*, Vol. 54, (1), November 2000, pp. 57–90.
 - [23] Meurant, G., “Multitasking the Conjugate Gradient method on the CRAY X-MP/48”, *Parallel Computing*, Vol. 5, (3), July 1987, pp. 267–280.
 - [24] Saad, Y., “Krylov Subspace Methods on Supercomputers,” *SIAM Journal on Scientific and Statistical Computing*, Vol. 10, (6), November 1989, pp. 1200–1232.
 - [25] Daniel, J. W., Gragg, W. B., Kaufman, L. and Stewart, G. W., “Reorthogonalization and Stable Algorithms for Updating the Gram-Schmidt QR Factorization,” *Mathematics of Computation*, Vol. 30, (136), October 1976, pp. 772–795.
 - [26] Giraud, L., Langou, J. and Rozioznik, M., “The Loss of Orthogonality in the Gram-Schmidt Orthogonalization Process,” *Computers and Mathematics with Applications*, Vol. 50, (7), October 2005, pp. 1069–1075.